CHICAGO JOURNAL OF THEORETICAL COMPUTER SCIENCE 2016, Article 10, pages 1–16
http://cjtcs.cs.uchicago.edu/

Circuit Complexity of Powering in Fields of Odd Characteristic

Arkadev Chattopadhyay[*][†] Frederic Green[†] Howard Straubing

Received November 10, 2014; Revised March 28, 2016, and in final form July 28, 2016; Published August 3, 2016

Abstract: By a careful analysis of the proofs of Kopparty in fields of characteristic 2, we show that the problem of powering in $\mathbb{F}_{p^n}$ requires $\mathsf{ACC(p)}$ circuits of exponential size (in n), for any fixed prime $p > 2$. Similar bounds hold for quadratic residuosity. As a corollary, we obtain non-trivial bounds for exponential sums that express the correlation between the quadratic character of $\mathbb{F}_{p^n}$ and n-variate polynomials over $\mathbb{F}_p$ of degree up to n^{ε} for some $0 < \varepsilon < 1$.

1 Introduction

In this paper, we study the problem of powering in finite fields of odd characteristic. The same problem was studied recently by Kopparty [3] in fields of characteristic 2. For example, in [3] it is shown that, for appropriate values of n, finding the cube root of an element of $\mathbb{F}_{2^n}$ requires exponential size circuits of bounded depth with AND, OR, and parity gates (so-called $\mathsf{ACC(2)}$ or $\mathsf{AC}(\oplus)$ circuits), and that similar lower bounds hold for computing cubic residuosity. Here we show, for any odd prime p, that finding the square root of an element of $\mathbb{F}_{p^n}$ (under the promise that one exists) requires exponential size circuits consisting of AND, OR, and Mod_p gates (known as $\mathsf{ACC(p)}$-type circuits). This follows easily from the fact (also proved here) that quadratic residuosity requires exponential size for all n. Similar results hold for more general powering (e.g. cube roots for all primes $p > 3$) as well as other residue symbols.

[*] Partially supported by a Ramanujan fellowship of the DST, India.

[†] Work completed in part at Schloß Dagstuhl Seminar 14391, Algebra in Computational Complexity, 9/21/2014–9/26/2014.

Key words and phrases: circuit complexity, lower bounds, finite fields, exponential sums

DOI: 10.4086/cjtcs.2016.010

Perhaps the most intriguing corollaries in [3] are upper bounds for a certain class of character sums. These bounds, which essentially show that various resdiuosity symbols do not correlate well with low-degree polynomials, are interesting because they are complete (with the variable of summation running over the entire field), and hold for polynomials of degree up to n^{ε}, for $0 < \varepsilon < 1$. Classical techniques (e.g., those based on Weyl differencing and its generalizations; see [1]) for this type of sum can give bounds this sharp only for degree less than $\log n$. Recently, we received a communication [4] that, prior to and independently of our work, Kopparty had succeeded in generalizing his original characteristic 2 character sum bounds to arbitrary characteristic. However, he obtained these results directly, without using lower bounds for $\mathsf{ACC(p)}$ circuits. What distinguishes our work from his is that we show that the entire apparatus of his original proof works for any characteristic, in particular including the $\mathsf{ACC(p)}$ circuit lower bounds.

2 Encoding Field Elements and $\mathbb{F}_p$ Approximation of $\mathsf{ACC(p)}$ circuits

Our lower bounds rely on encoding field elements as strings of bits. This in turn requires restricting the set of boolean inputs to a proper subset of all 2^m input settings of an $\mathsf{ACC(p)}$ circuit with m input bits. Hence some care must be taken to ensure that the circuits are not handicapped in an unnatural way, so that under one encoding they would be weak, whereas under another they would be capable of powering. It is important to point out that our results apply *for any reasonable encoding of the field elements*. It is therefore essential to discuss briefly what we mean by “reasonable encoding."

Given that p is constant as m varies, it is intuitively clear that a reasonable encoding should only entail a fixed (i.e., m-independent) number of boolean inputs. Thus, we define an encoding to be *reasonable* if it is a one-to-one function $e : \mathbb{F}_p \to \{0,1\}^r$ for some $r \in \mathbb{N}$ independent of m. The circuit would effectively accept field elements as inputs by breaking the boolean inputs up into blocks of size r (in an arbitrary way, but using the same r for all m) such that the bits in each block take on values in the image of e. Since any such function e can be represented by a fixed-degree polynomial over $\mathbb{F}_p$, the constructions below that produce polynomial representations of circuits work for any encoding, and hence the approximation results (and subsequent lower bounds) apply to $\mathsf{ACC(p)}$ circuits under any encoding scheme.

We return to this issue more formally at the end of this section, where we provide further justification for the notion of “reasonable encoding." There we demonstrate that, for reasonable encodings, the circuit sizes in our approximation results (Lemma 2.1 and Corollary 2.2) are *polynomially related*, and the degrees of approximating polynomials are determined to within a *constant factor*. This is sufficient for our results.

While the subsequent development does not depend in any essential way on the choice of a reasonable encoding, for concreteness we now describe one such simple encoding scheme, which, the reader may assume, applies for the rest of the paper. We view the elements of $\mathbb{F}_p$ as $\{0,1,\ldots,p-1\}$, and we define a function $e : \mathbb{F}_p \to \{0,1\}^{p-1}$ by setting $e(x) = 0^{p-1-x}1^x$ for any $x \in \mathbb{F}_p$. We can realize the mapping e as a $(p-1)$-tuple of polynomials $e_i : \mathbb{F}_p \to \mathbb{F}_p$, $0 \le i \le p-2$. Simply set $e_i(x) = (\prod_{j=0}^{i}(x-j))^{p-1}$ for $0 \le i \le p-2$. Each e_i is a polynomial of degree $O((p-1)^2) = O(1)$. Furthermore, in $\mathbb{F}_p$, $e_i(x) = 1$ if $i < x$ and is 0 otherwise. Hence, regarding 0 and 1 as elements of $\mathbb{F}_p$, $e(x)$ is the bit string $e_{p-2}(x)e_{p-3}(x)\cdots e_0(x)$.

Consider any boolean circuit C. Suppose that the number of boolean inputs to C is m, labelled $y_1,\ldots,y_m$. Unless noted otherwise, we will assume throughout that $(p-1)|m$, so that we may write

$m = n(p-1)$ for some $n \in \mathbb{N}$. Arrange the inputs into blocks of size $p-1$. Thus the first block is $y_1, \ldots, y_{p-1}$, and the i^{th} is $y_{i(p-1)+1}, \ldots, y_{(i+1)(p-1)}$, for $1 \leq i \leq n$. The $\mathbb{F}_p$ *input block restriction of* C or, for short, the *block restriction of* C, is the set of inputs of the form $\prod_{i=1}^{n}(0^{k_i}1^{p-1-k_i})$, where each $k_i \in \{0, \ldots, p-1\}$. Thus for a particular input setting, in block i, $y_j = 0$ for $i(p-1) < j \leq i(p-1)+k_i$ and $y_j = 1$ for $i(p-1)+k_i < j \leq (i+1)(p-1)$. Note that there are p^n input settings of this form, and that each of the n blocks of $p-1$ inputs to C can be understood as an element of $\mathbb{F}_p$ according to the encoding scheme given above.

In the obvious way, we now define a version of C, denoted $C_{\mathbb{F}_p}$, which behaves exactly like C under the block restriction, but which takes as input n elements of $\mathbb{F}_p$, denoted $x_1, \ldots, x_n$. $C_{\mathbb{F}_p}(x_1, \ldots, x_n)$ is defined as the composition of $C(y_1, \ldots, y_m)$ with the polynomials e_j, such that in block i, we replace the j^{th} input $y_{i(p-1)+j}$ to C by $e_{p-1-j}(x_i)$. (Note that, although $C_{\mathbb{F}_p}$ takes *bona fide* elements of $\mathbb{F}_p$ as inputs, it is not an arithmetic circuit over $\mathbb{F}_p$. It is an $\mathsf{ACC(p)}$ circuit “fooled into thinking that it is taking $\mathbb{F}_p$ inputs." The *output* of $C_{\mathbb{F}_p}$ is, at this point, still boolean.)

A key element of the proof is to approximate C, via $C_{\mathbb{F}_p}$, by a low-degree polynomial $t : \mathbb{F}_p^n \to \{0,1\}$, as in the method of Razborov and Smolensky [6, 7]. Because, under the block-restriction, the distribution of the inputs is not uniform over $\{0,1\}^m$, this is most easily done via the well-known technique of probabilistic polynomials [5]. A *probabilistic polynomial* q over $\mathbb{F}_p$ is a multivariate polynomial with two disjoint sets of variables: ordinary variables $x_1, \ldots, x_n$, and probabilistic variables $b_1, \ldots, b_r$. To evaluate q on a bit string $\mathbf{x}$ of length n, we choose a $\mathbf{b}$ of bits uniformly and at random from $\{0,1\}^r$ and use the value $q(\mathbf{x}, \mathbf{b})$, treating the boolean inputs as elements of $\mathbb{F}_p$. We say that q represents the boolean function $f : \{0,1\}^n \to \{0,1\}$ with error

$$\mathsf{Pr}_{\mathbf{b} \in \{0,1\}^r}[f(\mathbf{x}) \neq q(\mathbf{x}, \mathbf{b})].$$

Lemma 2.1. *Let C be an $\mathsf{ACC(p)}$ circuit of depth d and size p^{n^δ}, where $\delta < 1/(16d)$, subject to the $\mathbb{F}_p$ input block restriction. Then there exists a polynomial $t : \mathbb{F}_p^n \to \{0,1\}$ of degree $O(n^{1/4})$ such that $C_{\mathbb{F}_p}(x_1, \ldots, x_n) = t(x_1, \ldots, x_n)$ on a fraction $1 - p^{-n^{1/8d}}$ of the p^n input settings.*

Proof. Denote the size of C by $s = O(p^{n^\delta})$. Let $\varepsilon = p^{-n^{1/8d}}$. In [5] it is shown that each AND and OR gate in C can be represented as a probabilistic polynomial over $\mathbb{F}_p$, of degree $O((\log s)(\log(s/\varepsilon)))$, error no more than ε/s and with $O(\log^2 s \log(s/\varepsilon))$ probabilistic variables. Use the same probabilistic variables for all AND and OR gates, and compose the polynomials to obtain polynomial t', whose variables are the block-restricted inputs to C. The Mod$_p$ gates are implemented by raising to the power $p-1$ and appealing to Fermat's little theorem, and hence only multiply the degree of t' by a constant independent of n. Thus since the depth is d, the degree of t' is $O((\log s)^d)(\log(s/\varepsilon))^d) = O(n^{2d\delta+1/8}) = O(n^{1/4})$, and the probability that some AND or OR gate errs is bounded above by $s \cdot \varepsilon/s = \varepsilon = p^{-n^{1/8d}}$. By an averaging argument, there exists a setting of the probabilistic variables such that the polynomial t' agrees with C on a fraction $1-\varepsilon$ of the block-restricted input settings. Composing t' with the $O(1)$-degree polynomials e_i defined above yields a polynomial t, also of degree $O(n^{1/4})$, over $\mathbb{F}_p$ that agrees with $C_{\mathbb{F}_p}$ on a fraction $1-\varepsilon$ of the $\mathbb{F}_p$-valued inputs. □

It is quite straightforward to generalize the circuit C to have multiple outputs, and hence via a suitable block restriction of the outputs, to represent functions from $\mathbb{F}_p^n \to \mathbb{F}_p^n$, in which blocks of $p-1$ boolean outputs represent field elements. Using a $(p-1)$-bit unary encoding as in the inputs, the $\mathbb{F}_p$ value would simply be the sum of the bits in a block, which varies from 0 through $p-1$. Once again we denote the multi-output boolean $\mathsf{ACC(p)}$ circuit realizing the encoding by C, and the equivalent $\mathbb{F}_p^n$-valued circuit by $C_{\mathbb{F}_p}$. The following is then an immediate corollary of Lemma 2.2.

Corollary 2.2. *Let C be an $\mathsf{ACC(p)}$ circuit of depth d and size p^{n^δ}, where $\delta < 1/(16d)$, subject to the $\mathbb{F}_p$ input and output block restriction. Then there exists an n-tuple of polynomials $(t_1,\ldots,t_n)$, with $t_i : \mathbb{F}_p^n \to \mathbb{F}_p$ for each $1 \le i \le n$, where each t_i has degree $O(n^{1/4})$, such that $C_{\mathbb{F}_p}(x_1,\ldots,x_n) = (t_1(x_1,\ldots,x_n),\ldots,t_n(x_1,\ldots,x_n))$ on a fraction $1-p^{-n^{1/8d}}$ of the $\mathbb{F}_p^n$ input settings.*

We return to the issue mentioned at the beginning of this section regarding encodings. Observe that, in the statement and proofs of Lemma 2.1 and Corollary 2.2, no explicit reference is made to the particular encoding. In fact the encoding has no significant effect on the proofs or results, and we formalize this fact in the following proposition.

Proposition 2.3. *For any reasonable encoding, the degrees of the polynomials t and t_i constructed in Lemma 2.1 and Corollary 2.2 differ by at most a constant factor. The sizes of the resulting circuits are polynomially related.*

Proof. The polynomial t that is constructed in Lemma 2.1 is the composition of a degree $O(n^{1/4})$ probabilistic polynomial for the *boolean* circuit C with the encoding polynomials e_i. The e_i have degree $O(1)$, and that constant factor is the only way in which the encoding affects the degree. A different reasonable encoding would only multiply the degree of t by $O(1)$. In Corollary 2.2 the encoding has no additional affect on the degrees of the polynomials t_i, since the $\mathbb{F}_p$-valued outputs are linear functions (sums) of the boolean outputs. Thus the encoding scheme has no affect on the degrees of the polynomials in our approximations, outside of constant factors, which are ignored.

Our circuit sizes are all expressed as a function of n, the number of field elements, which equals m/r for an encoding of block size r. Let the block size of an encoding scheme e be denoted by r_e. Thus for the encoding scheme defined in this section, $r_e = p-1$. We now show that, for reasonable encodings, the relevant circuit sizes are polynomially related.

Our lower bounds are all of the form $s(n) = p^{n^\delta}$. Consider any encoding scheme e' using block size $r_{e'}$. Then

$$s(m/r_{e'}) = p^{(m/r_{e'})^\delta} = p^{(\frac{m}{p-1})^\delta(\frac{p-1}{r_{e'}})^\delta} = \left(p^{(\frac{m}{p-1})^\delta}\right)^{(\frac{p-1}{r_{e'}})^\delta} = (s(m/(p-1)))^{(\frac{p-1}{r_{e'}})^\delta}.$$

For any reasonable encoding $r_{e'}$, $\frac{p-1}{r_{e'}} = \Theta(1)$ and hence $s(m/r_{e'}) = \Theta(s(m/(p-1))^c)$ for some constant c. □

Note that reasonable encodings are not only sufficient but also necessary for the sizes to be polynomially related: By the proof above, if the two sizes $p^{(m/r_{e'})^\delta}$ and $p^{(m/(p-1))^\delta}$ are polynomially related, then e' is a reasonable encoding scheme. An analogous assertion applies to the degrees of the approximating polynomials. We take this as further justification of our interpretation of e as a reasonable encoding.

3 Versatility of Powering

3.1 Pseudorandomness

We define the mod-p weight $\mathrm{wt}_{\mathbb{F}_p}(j)$ of $j \in \mathbb{Z}_{p^n-1}$ to be the sum of the digits in the radix p expansion of j. Thus $\mathrm{wt}_{\mathbb{F}_p}(j)$ generalizes the Hamming weight of j to an odd prime radix. We then set

$$M_d = \{j : 0 \leq j < p^n - 1, \mathrm{wt}_{\mathbb{F}_p}(j) \leq d\}.$$

For technical reasons we view M_d as a subset of $\mathbb{Z}_{p^n-1}$ rather than $\mathbb{Z}_{p^n}$.

Note that $(p-1)n/2$ is close to half the maximum base-p weight of an element of $\mathbb{Z}_{p^n-1}$. (It is not exactly half, because we exclude the radix p representation of $p^n - 1$, which would be the only element of weight $(p-1)n$.) If $\alpha \in \mathbb{Z}_{p^n-1}$, each of $M_{(p-1)n/2}$ and $\alpha + M_{(p-1)n/2}$ contain approximately half the elements of $\mathbb{Z}_{p^n-1}$: The set of integers j with weight strictly less than $(p-1)n/2$ has one more element than the set of j with weight strictly greater than $(p-1)n/2$, and those with weight exactly $(p-1)n/2$ form a vanishingly small fraction of the total as n increases.

In the case of $p = 2$, a key technical result of Kopparty's study of the complexity of powering concerned the behavior of $\alpha \in \mathbb{Z}_{2^n-1}$ with respect to the set $M_{n/2}$ (i.e., the set of elements of $\mathbb{Z}_{2^n-1}$ of Hamming weight at most $n/2$). In particular, for certain choices of α, the set of translates $M_{n/2} + \alpha$ of the elements of $M_{n/2}$, and $M_{n/2}$ itself, cannot have a very large intersection. The same phenomenon applies to odd primes p.

The next theorem formalizes this. It shows that if $\alpha \in \mathbb{Z}_{p^n-1}$ is chosen properly, then r and $r + \alpha$, where r is a randomly chosen integer in $\mathbb{Z}_{p^n-1}$, behave like a pair of independent random choices from $\mathbb{Z}_{p^n-1}$. As a result, the overlap between $M_{(p-1)n/2}$ and $M_{(p-1)n/2} + \alpha$ cannot be very large; more precisely, $M_{(p-1)n/2} \cup (\alpha + M_{(p-1)n/2})$ contains significantly more than half the elements of $\mathbb{Z}_{p^n-1}$.

Theorem 3.1. *(Pseudorandomness of Periodic α) Let $D = \{0, 1, \ldots, p-1\}$. Let t, t' be constants. Let $\sigma_0 \in D^t$ be a string of digits base p, with $\sigma_0 \neq 0^t, (p-1)^t$, and $\sigma' \in D^{t'}$.*

Let $\sigma \in D^n$ be a string of base p digits of the form $\sigma' \sigma_0^{\ell}$, with $n = t' + \ell t$. Let $\alpha \in \mathbb{Z}_{p^n-1}$ be the integer whose base p representation is σ.

Then there exists $\varepsilon > 0$ such that for all sufficiently large n, depending only on σ_0, σ', we have,

$$|M_{n(p-1)/2} \cup (\alpha + M_{n(p-1)/2})| \geq (\frac{1}{2} + \varepsilon) p^n.$$

The proof of this theorem is essentially identical to that of the base 2 version in [3]. All changes are due to the fact that we are working in base p rather than base 2. However, given the centrality of this theorem to our main results, a sketch of the proof (closely following the one given in [3]) is included in Appendix A.

3.2 Versatility

We use the pseudorandomness result of the preceding subsection to show that powering is "versatile," in the sense that when the powering function is available, computing any function in $\mathbb{F}_{p^n}$ effectively reduces to evaluating multivariate polynomials of "low" (in this case, $\leq n(p-1)/2$) degree over $\mathbb{F}_p$, or in a space of “low" (in this case, $\leq (1/2-\varepsilon)p^n$) dimension. Our treatment departs from Kopparty’s in some technical details. In particular, in the statement of Theorem 3.6 below, the upper bound on the dimension of V is smaller by 2, which eliminates the need for an extra case in the proof. This is achieved simply by restricting the domain to $\mathbb{F}^*_{p^n}$ rather than $\mathbb{F}_{p^n}$, but this in turn entails some subtleties, and we therefore provide more detail on this result.

Let p be prime, $n > 0$, and let $\mathbb{F}_{p^n}$ be the field with p^n elements. This is an n-dimensional vector space over $\mathbb{F}_p$, and thus there is an $\mathbb{F}_p$-linear isomorphism η between $\mathbb{F}_{p^n}$ and $\mathbb{F}^n_p$. The argument works for any such η (there is one point at which this is not obvious, and there we explain why it works nonetheless). We remove the zero vector from both these representations of the vector space, and thus deal with $\mathbb{F}^*_{p^n}$, the group of units of the field, and with the set which we denote, by an abuse of notation, $(\mathbb{F}^n_p)^*$, the set of nonzero n-tuples of elements from $\mathbb{F}_p$.

We make a few observations about functions whose domains are these sets.

Lemma 3.2. *Every function*

$$f : \mathbb{F}^*_{p^n} \to \mathbb{F}_{p^n}$$

has a unique representation as a univariate polynomial

$$\sum_{j=0}^{p^n-2} a_j x^j,$$

where $a_j \in \mathbb{F}_{p^n}$ for each j.

Proof. This is simply Lagrange interpolation, which gives a unique polynomial of degree strictly less than $p^n - 1$ when $p^n - 1$ function values are specified. □

We denote by $\deg_{\mathbb{F}_{p^n}}(f)$ the degree of this unique polynomial representation.

Because of the isomorphism η each such function f can also be viewed as a function $\eta \circ f \circ \eta^{-1} : (\mathbb{F}^n_p)^* \to \mathbb{F}^n_p$, which we still write as,

$$f : (\mathbb{F}^n_p)^* \to \mathbb{F}^n_p.$$

Lemma 3.3. *With respect to this representation, every function has a unique representation as an n-tuple*

$$(q_1(y_1,\dots,y_n),\dots,q_n(y_1,\dots,y_n))$$

of multivariate polynomials over $\mathbb{F}_p$, where each monomial has the form

$$y_1^{e_1}\cdots y_n^{e_n},$$

with $0 \leq e_i < p$ for all i, and $e_i < p-1$ for at least one i.

Proof. Without the restriction to nonzero domain elements and the absence of the term of maximal degree, this fact is well known. We need to see how to obtain the modification stated here. Let $a_1,\dots,a_n \in \mathbb{F}_p$. The polynomial

$$\prod_{i=1}^{n}((y_i-a_i)^{p-1}-1)$$

is 0 at every point of $\mathbb{F}_p^n$ except $(a_1,\cdots,a_n)$, at which it has value $(-1)^n$. Thus every function from $\mathbb{F}_p^n$ into $\mathbb{F}_p$ can be written as a linear combination of monomials in which the degree of each y_i is strictly less than p. As the set of such functions is a p^n-dimensional vector space over $\mathbb{F}_p$, and there are p^n such monomials (counting the degree 0 monomial 1), the set of these monomials is linearly independent over $\mathbb{F}_p$.

We call the monomials other than $y_1^{p-1}\cdots y_n^{p-1}$ *submaximal*. We claim that the set of submaximal monomials is still linearly independent when considered as functions with domain $(\mathbb{F}_p^n)^*$. If not, there would be a nontrivial linear combination of these that is identically zero on $(\mathbb{F}_p^n)^*$. The resulting polynomial r has degree strictly less than $(p-1)\cdot n$. We have $r(0,\ldots,0)=c$, for some $c\in\mathbb{F}_p$, and $r(a_1,\ldots,a_n)=0$ elsewhere. Now consider the polynomial

$$s(y_1,\ldots,y_n)=(-1)^n c\cdot\prod_{i=1}^{n}(y_i^{p-1}-1).$$

We have $s(a_1,\ldots,a_n)=r(a_1,\ldots,a_n)$ for all $a_1,\ldots,a_n\in\mathbb{F}_p$, but the two polynomials have different degrees. This contradicts linear independence of the complete set of p^n monomials, and thus establishes linear independence of the submaximal monomials with domain restricted to $(\mathbb{F}_p^n)^*$. As the set of functions from $(\mathbb{F}_p^n)^*$ into $\mathbb{F}_p$ is a space of dimension p^n-1, the submaximal monomials form a basis for this space, which gives the desired conclusion. □

Given such a multivariate polynomial representation $(q_1,\ldots,q_n)$ of f, we define the *$\mathbb{F}_p$-degree of f* as,

$$\deg_{\mathbb{F}_p}(f)=\max_i \deg(q_i).$$

This appears to depend upon the isomorphism $\eta:\mathbb{F}_{p^n}\to\mathbb{F}_p^n$. But note that if we have two such isomorphisms, they can be composed to form a linear automorphism of $\mathbb{F}_p^n$. When we translate the polynomial representation of a function via such an automorphism, we merely introduce a linear change of variables, and thus the degree is not affected.

There is a tight connection between these two different interpretations of the degree of a function. Let $0\le i<p^n-1$. Observe that the maximum possible weight $\mathrm{wt}_{\mathbb{F}_\mathrm{p}}(i)$ is $(p-1)n-1$, since we exclude $i=p^n-1$. The minimum nonzero weight is 1, attained only when i is an integer power of p. The proof of the following can be found in Kopparty's thesis [2].

Theorem 3.4. *Let $0\le d\le n(p-1)$, and let $f:\mathbb{F}_{p^n}^*\to\mathbb{F}_{p^n}$. Then $\deg_{\mathbb{F}_p}(f)\le d$ if and only if the univariate $\mathbb{F}_{p^n}$ representation of f is a linear combination of monomials x^i with $\mathrm{wt}_{\mathbb{F}_\mathrm{p}}(i)\le d$.*

Observe that a univariate polynomial f over $\mathbb{F}_{p^n}$ of low $\mathbb{F}_p$-degree can have a degree that is quite high. For example, the monomial $x^{p^{n-1}}$ has $\mathbb{F}_p$-degree $\mathrm{wt}_{\mathbb{F}_\mathrm{p}}(p^{n-1})=1$. However, low $\mathbb{F}_p$-degree, which corresponds to a sparse univariate polynomial over $\mathbb{F}_{p^n}$, is sufficient for our purposes.

From Corollary 2.2, the following is immediate.

Corollary 3.5. *Let C be an $\mathsf{ACC(p)}$ circuit of depth d and size p^{n^δ}, where $\delta<1/(16d)$, subject to the $\mathbb{F}_p$ input and output block restriction. Then there exists a univariate polynomial $t:\mathbb{F}_{p^n}\to\mathbb{F}_{p^n}$ of $\mathbb{F}_p$-degree $O(n^{1/4})$ such that $C_{\mathbb{F}_p}(x)=t(x)$ on a fraction $1-p^{-n^{1/8d}}$ of the $\mathbb{F}_{p^n}$ input settings.*

Hence the circuit lower bounds we are interested in reduce to establishing the computational limitations of low $\mathbb{F}_p$-degree (i.e., $O(n^{1/4})$) polynomials.

Now we adapt the result of [3] that the powering function $\Lambda : x \mapsto x^\alpha$ on $\mathbb{F}^*_{p^n}$ is *versatile* for suitable α. Roughly speaking, this means that any function from $\mathbb{F}^*_{p^n}$ into $\mathbb{F}_{p^n}$ can be expressed in terms of Λ in an "easy way."

Theorem 3.6. *(Versatility of Powering) Let $0 < \varepsilon < 1/2$, and $\alpha \in \mathbb{Z}_{p^n-1}$ be such that,*

$$|M_{n(p-1)/2} \cup (\alpha + M_{n(p-1)/2})| \geq (\frac{1}{2} + \varepsilon)p^n.$$

Let $\Lambda : \mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ be defined by $\Lambda(x) = x^\alpha$.

Then there exists an $\mathbb{F}_{p^n}$-linear space V of functions from $\mathbb{F}^*_{p^n}$ to $\mathbb{F}_{p^n}$ such that the following holds. Let $f : \mathbb{F}^*_{p^n} \to \mathbb{F}_{p^n}$. Then there are univariate polynomials $g, h, e : \mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ such that,

- $\deg_{\mathbb{F}_p}(g), \deg_{\mathbb{F}_p}(h) \leq (p-1)n/2$,
- $\dim(V) \leq \left(\frac{1}{2} - \varepsilon\right) \cdot p^n$,
- *for every $x \in \mathbb{F}^*_{p^n}$,*
$$f(x) = g(x) + \Lambda(x) \cdot h(x) + e(x), \tag{3.1}$$
where $e \in V$.

Proof. It is enough to specify a V such that for each $i < p^n - 1$, there exist polynomials g and h of the required degrees for which Eq. (3.1) holds with $f(x) = x^i$. We obtain the stated result by summing over all the relevant monomials x^i, using the univariate polynomial representation of f.

We let V be the subspace spanned by the monomials

$$\{x^i : i \notin M_{(p-1)n/2} \cup (\alpha + M_{(p-1)n/2})\}$$

Observe that

$$\begin{aligned} \dim(V) &= p^n - 1 - |M_{(p-1)n/2} \cup (\alpha + M_{(p-1)n/2}| \\ &\leq p^n - (\frac{1}{2} + \varepsilon) \cdot p^n \\ &= (\frac{1}{2} - \varepsilon) \cdot p^n. \end{aligned}$$

If x^i is one of these monomials, then we have $x^i = h + g \cdot \Lambda + e$, where h and g are identically zero, and $e(x) = x^i$. If $i \in M_{(p-1)n/2}$, then we have $x^i = g + h \cdot \Lambda + e$, by taking h and e to be identically zero and $g(x) = x^i$. The bound on the $\mathbb{F}_p$-degree of g follows from Theorem 3.4. The final case is if $i \in \alpha + M_{(p-1)n/2}$, in which case we have $i = (\alpha + j) \bmod (p^n - 1)$ for some $j \in M_{(p-1)n/2}$. We set g and e to be identically 0 and $h(x) = x^j$. Again by Theorem 3.4, we get $\deg_{\mathbb{F}_p} h \leq (p-1)n/2$. Note that for all nonzero x,

$$x^i = x^{(\alpha+j) \bmod (p^n-1)} = x^{\alpha+j} = g(x) + h(x) \cdot \Lambda(x) + e(x).$$

Let us see why we can remove the reduction modulo $p^n - 1$ in the above sequence of equations. Either $\alpha + j = (\alpha + j) \bmod (p^n - 1)$, in which case there is nothing to prove, or $\alpha + j = p^n - 1 + (\alpha + j) \bmod (p^n - 1)$, in which case

$$x^{\alpha+j} = x^{(\alpha+j) \bmod (p^n-1)} \cdot x^{p^n-1} = x^{\alpha+j},$$

because $x^{p^n-1} = 1$ for all nonzero $x \in \mathbb{F}_{p^n}$.

□

Note now that Theorem 3.1 gives a class of powers α that satisfy the requirement for $|M_{n(p-1)/2} \cup (\alpha + M_{n(p-1)/2})|$ in the hypothesis of the above theorem. For these values of α, the versatility of x^α yields hardness results as described in the next two sections.

4 Hardness of Quadratic Residuosity and Powering

Let $q > 1$. Given an input $x \in \mathbb{F}_{p^n}$ we would like to determine if there exists $y \in \mathbb{F}_{p^n}$ such that $x = y^q$. This is the q^{th} *power residuosity problem.* In [3] it is shown that computing cubic residuosity (and indeed q^{th} residuosity for any odd q) is hard in the characteristic 2 setting. Here we show that in the characteristic p setting, *quadratic* residuosity is hard. We also find that q^{th} residuosity is hard for any q relatively prime to p. We denote the *quadratic character* in $\mathbb{F}_{p^n}$ by $\chi(x) = x^\alpha$ where $\alpha = (p^n - 1)/2$. Then $\chi(x) = 1$ if and only if x is a quadratic residue in $\mathbb{F}_{p^n}$. The base p expansion of α consists of n digits, each equal to $(p-1)/2$, so α has period 1.

To illustrate how to generalize Kopparty's proof in this setting, we directly prove a weaker version of the hardness of quadratic residuosity. The exact analog of this is not present in [3] (a proof is outlined in section 2.2 of the paper), but is an easy corollary of stronger results (in particular, Theorem 3 of [3]). The essence of this proof is the Razborov/Smolensky dimensionality argument, coupled with the $\mathbb{F}_p^n$–$\mathbb{F}_{p^n}$ polynomial correspondence.

Theorem 4.1. *Let C be an* $\mathsf{ACC(p)}$ *circuit of depth d, under the $\mathbb{F}_p$ block restriction, such that $C_{\mathbb{F}_p}$ computes a function from $\mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$, let $\alpha = (p^n - 1)/2$, and let $\chi(x) = x^\alpha$ be the quadratic character in $\mathbb{F}_{p^n}$. Then if the size of C is at most p^{n^δ} where $\delta < 1/(16d)$, for sufficiently large n, we have:*

$$\Pr_{x \in \mathbb{F}_{p^n}} [C_{\mathbb{F}_p}(x) = \chi(x)] \leq 1 - \varepsilon,$$

where $\varepsilon < 1$.

Proof. By Corollary 3.5, there exists a polynomial $t : \mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ of degree $O(n^{1/4})$ such that on a fraction $1 - p^{-n^{1/8d}}$ of all inputs, $t(x) = C_{\mathbb{F}_p}(x)$. Let $S = \{x \in \mathbb{F}_{p^n} | t(x) = C_{\mathbb{F}_p}(x)\}$ denote the set of inputs on which there is agreement. Thus $|S| \geq (1 - p^{-n^{1/8d}})p^n$. Let A denote the set of inputs for which the circuit agrees with quadratic residuosity, i.e., $C_{\mathbb{F}_p}(x) = \chi(x)$.

In order to invoke Theorem 3.6, we must show that α has the requisite properties given in Theorem 3.1. As explained above, the base p representation of $\alpha = (p^n - 1)/2$ just consists of n repetitions of the digit $(p-1)/2$, i.e., it has period 1. Thus, in the notation of Theorem 3.1, we have $t = 1$, $\sigma_0 = (p-1)/2$

(which equals neither 0 nor $p-1$), $t' = 0$ and $\sigma' = \lambda$ (where λ denotes the empty string), and $\ell = n$. Thus Theorem 3.1 applies, for some $\varepsilon' > 0$ the inequality $|M_{n(p-1)/2} \cup (\alpha + M_{n(p-1)/2})| \geq (\frac{1}{2} + \varepsilon')p^n$ holds, so we can apply Theorem 3.6. That is, for any function $f : \mathbb{F}_{p^n}^* \to \mathbb{F}_{p^n}$, we may write $f(x) = g(x) + \chi(x) \cdot h(x) + e(x)$, where $\deg_{\mathbb{F}_p}(g), \deg_{\mathbb{F}_p}(h) \leq n(p-1)/2$ and $e \in V$ where $\dim(V) \leq (\frac{1}{2} - \varepsilon')p^n$.

Now consider the functions $f|_{S \cap A}$ restricted to $S \cap A$. If $x \in S \cap A$, we know $t(x) = \chi(x)$, so that we may write $f|_{S \cap A}(x) = g(x) + t(x)h(x) + e(x)$. Note that $\deg_{\mathbb{F}_p}(t \cdot h) \leq n(p-1)/2 + O(n^{1/4})$. Given the upper bound on the dimension of V, we can then place an upper bound on the dimension of the space of functions over $S \cap A$. The polynomials of the form $g + t \cdot h$ are spanned by monomials of degree up to $n(p-1)/2 + O(n^{1/4})$, and thus this space has dimension at most $(\frac{1}{2} + O(n^{1/4}/\sqrt{n}))p^n = (\frac{1}{2} + O(1/n^{1/4}))p^n$. Including the dimension of V, the space of functions spanned by polynomials of the form $g(x) + t(x)h(x) + e(x)$ has dimension $(1 + O(n^{-1/4}) - \varepsilon')p^n$.

On the other hand, the space of functions $f|_{S \cap A}$ contains *all* functions on the set $S \cap A$, and thus has dimension $|S \cap A|$. Thus $|S \cap A| \leq (1 + O(n^{-1/4}) - \varepsilon')p^n$.

Since $|S| \geq (1 - p^{-n^{1/8d}})p^n$, we conclude that $|\overline{S}| \leq p^{n - n^{1/8d}}$. This implies that $|A| \leq |S \cap A| + |\overline{S}| \leq (1 + O(n^{-1/4}) + p^{-n^{1/8d}} - \varepsilon')p^n \leq (1 - \varepsilon'/2)p^n$ for sufficiently large n. Taking $\varepsilon = \varepsilon'/2$ and noting that $|A| = p^n \Pr_{x \in \mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = \chi(x)]$ proves the theorem. □

With minor modifications, this can be generalized to account for any power that has the appropriate pseudorandomness properties. We omit the proof, as it follows closely the argument given above and is essentially the same as the analog given in Kopparty's paper (Theorem 1 in [3]).

Theorem 4.2. *Let C be an $\mathsf{ACC(p)}$ circuit of depth d, under the $\mathbb{F}_p$ block restriction, such that $C_{\mathbb{F}_p}$ computes a function from $\mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$.*

Let $a, b, q \in \mathbb{Z}$ be constants, where $q > 1$ is relatively prime to p, and $0 < |a|, |b| < q$. Let n be such that $\alpha = (ap^n + b)/q$ is an integer.

Let $\Lambda(x) = x^\alpha$. If the size of C is at most p^{n^δ} where $\delta < 1/(16d)$, then we have:

$$\Pr_{x \in \mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = \Lambda(x)] \leq 1 - \varepsilon,$$

for sufficiently large n, where $\varepsilon < 1$ depends only on q.

We next observe that the hardness of quadratic residuosity implies that finding square roots is hard, in the sense that given an $x \in \mathbb{F}_{p^n}$ with the promise that x has a square root, exponential size $\mathsf{ACC(p)}$ circuits are required to find that square root. Indeed, it is impossible for the circuit to output a square root (if it exists) for an appreciable fraction of the elements that have one. Analogous corollaries follow easily, in the same way, from Kopparty's result regarding cube (and higher) roots when not all field elements have such roots.

Corollary 4.3. *Let C be an $\mathsf{ACC(p)}$ circuit, under the $\mathbb{F}_p$ block restriction, such that $C_{\mathbb{F}_p}$ computes a function from $\mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$, and let χ be the quadratic character. Then if the size of C is at most p^{n^δ} where $\delta < 1/(16d)$, for sufficiently large n, we have:*

$$\Pr_{x \in \mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = x^{1/2} | \chi(x) = 1] \leq 1 - \varepsilon,$$

where $\varepsilon < 1$.

Proof. Using the same ε as in Theorem 4.1, suppose $\Pr_{x\in\mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = x^{1/2}|\chi(x) = 1] > 1-\varepsilon$. From $C_{\mathbb{F}_p}$ construct a new circuit $C'_{\mathbb{F}_p}$ which operates as follows: If x is 0, output 0. Otherwise, take the output of $C_{\mathbb{F}_p}$ and square it. If it is equal to x, output 1, otherwise -1. Thus, for any x, if $C_{\mathbb{F}_p}$ outputs the square root of x (which implies such a root exists), then $C'_{\mathbb{F}_p}(x) = 1 = \chi(x)$. On the other hand, for any x such that $\chi(x) = -1$, $C_{\mathbb{F}_p}$ will fail to output a square root, and hence $C'_{\mathbb{F}_p}(x) = -1 = \chi(x)$. That is, $C'_{\mathbb{F}_p} = \chi(x)$ for all x such that $\chi(x) \neq 1$ and hence $\Pr_x[C'_{F_p}(x) = -1 \wedge \chi(x) = -1] = \Pr_x[\chi(x) = -1]$. Now there are precisely $(p^n-1)/2$ values of x such that $\chi(x) = 1$ (respectively, $\chi(x) = -1$), and hence $\Pr_x[\chi(x) = 1] = \Pr_x[\chi(x) = -1] = \frac{1}{2}(1-p^{-n})$. Thus the inequality $\Pr_{x\in\mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = x^{1/2}|\chi(x) = 1] > 1-\varepsilon$ implies $\Pr_{x\in\mathbb{F}_{p^n}}[C_{\mathbb{F}_p}(x) = x^{1/2} \wedge \chi(x) = 1] > (1-\varepsilon)/2 + O(p^{-n})$. Hence,

$$\begin{aligned}
\Pr_x \quad & [C'_{F_p}(x) = \chi(x)] \\
= \quad & \Pr_x[C'_{F_p}(x) = -1 \wedge \chi(x) = -1] + \Pr_x[C'_{F_p}(x) = 1 \wedge \chi(x) = 1] + O(p^{-n}) \\
= \quad & \frac{1}{2} + \Pr_x[C'_{F_p}(x) = 1 \wedge \chi(x) = 1] + O(p^{-n}) \\
> \quad & \frac{1}{2} + (1-\varepsilon)/2 + O(p^{-n}) \\
= \quad & 1 - \varepsilon/2 + O(p^{-n}).
\end{aligned}$$

This contradicts Theorem 4.1. □

5 Stronger Bounds and Character Sums

We note here that Kopparty's proofs that self-reductions amplify the probabilities extend to characteristic p. The essential ingredients (approximation of the circuits by polynomials, and the fact that a fixed number of elements of $\mathbb{F}_{p^n}$ can be multiplied in $\mathsf{ACC(p)}$) have either been established in this paper or are known.

The following theorems are analogs of Theorems 2 and 3 in [3], respectively. The second generalizes Theorem 4.1 by both considering q^{th}-residuosity as well as showing sharper bounds on approximability.

Theorem 5.1. *Let $q > 1$ be odd (and constant), and choose n such that q and $p^n - 1$ are relatively prime. Let α be the inverse of q in $\mathbb{Z}_{p^n-1}$, and define the map $\Lambda : \mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ to be $\Lambda(x) = x^{\alpha}$. Let C be an $\mathsf{ACC(p)}$ block-restricted circuit of depth d and size $p^{n^{\delta}}$ where $\delta < 1/16d$. Then for sufficiently large n we have,*

$$\Pr_{x\in\mathbb{F}_{p^n}}[C(x) = \Lambda(x)] \leq 2^{-n^{\varepsilon}},$$

for some $\varepsilon < 1$ that depends on q and d.

Theorem 5.2. *Let $q > 1$ be an odd prime, and n sufficiently large that $q|(p^n-1)$.*

Let $\Lambda : \mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ be the q^{th} residue symbol in $\mathbb{F}_{p^n}$, i.e., $\Lambda(x) = x^{(p^n-1)/q}$.

If C is an $\mathsf{ACC(p)}$ *block-restricted circuit of depth d and size p^{n^δ} where $\delta < 1/(16d)$ then,*

$$\Pr_{x\in\mathbb{F}_{p^n}}[C(x) = \Lambda(x)] \le \frac{1}{q} + O(n^{-\varepsilon}),$$

where $\varepsilon < 1$ depends only on q and d.

Finally, we note that, as in [3], this leads to interesting bounds on character sums. Since the proof has some minor technical differences, we provide some detail for the quadratic character. Generalization to other multiplicative characters is straightforward.

Theorem 5.3. *There exists an $\varepsilon > 0$ such that the following holds. Let t be a polynomial in n variables over $\mathbb{F}_p$ of degree $\le n^\varepsilon$. Let $\eta : \mathbb{F}_{p^n} \to \mathbb{F}_p^n$ be an $\mathbb{F}_p$-linear isomorphism, $\chi : \mathbb{F}_{p^n} \to \mathbb{C}$ the quadratic character, and $\omega = e^{2\pi i/p}$ a primitive complex p^{th} root of unity. Then:*

$$\left| \frac{1}{p^n} \sum_{x\in\mathbb{F}_{p^n}} \chi(x)\omega^{t(\eta(x))} \right| \le n^{-\varepsilon}.$$

Proof. Denote

$$\sigma = \frac{1}{p^n} \sum_{x\in\mathbb{F}_{p^n}} \chi(x)\omega^{t(\eta(x))}.$$

Our goal, then, is to show that $|\sigma| \le n^{-\varepsilon}$. Define $\rho = \frac{\overline{\sigma}}{|\sigma|}$, so that $\rho\sigma = |\sigma|$. (Note in passing that $|\rho|^2 = 1$, so ρ is a point on the unit circle.) Multiplying the above equation by ρ, we obtain,

$$|\sigma| = \frac{1}{p^n} \sum_{x\in\mathbb{F}_{p^n}} \rho\chi(x)\omega^{t(\eta(x))},$$

and, taking the real part of both sides (and making use of the linearity of the operator $\Re(z) =$ the real part of z),

$$|\sigma| = \frac{1}{p^n} \sum_{x\in\mathbb{F}_{p^n}} \chi(x)\Re(\rho\omega^{t(\eta(x))}),$$

We now construct an $\mathsf{ACC(p)}$ circuit $C(x,r)$, under the block restriction, with $n(p-1)$ boolean inputs, and random bits r. $C(x,r)$ is designed to model a distribution related to $|\sigma|$. By composing $C(x,r)$ with the encoding polynomials e_i introduced earlier, at the same time we are constructing a corresponding $C_{\mathbb{F}_p}(x,r)$ with $\mathbb{F}_p$-valued inputs, and the same random bits r. $C_{\mathbb{F}_p}$ works as follows:

- Compute the polynomial t. This requires some justification. First note that $t \in \mathbb{F}_p[X_1,\dots,X_n]$ has total degree n^ε, and hence at most $\binom{n+n^\varepsilon}{n^\varepsilon} = 2^{n^{\varepsilon_1}}$ monomials, with $\varepsilon_1 < 1$. (The constant ε_1, and all the ε's used in this proof, are arbitrarily close to ε for sufficiently large n.) We claim that t can be computed by a $C_{\mathbb{F}_p}$-circuit of size $2^{n^{\varepsilon_2}}$, $\varepsilon_2 < 1$. To see this, first consider the computation of a single monomial of t, say $\prod_{i\in S} x_i^{d_i}$, where the exponents $d_i \in \mathbb{F}_p$, and by hypothesis, $\sum_{i\in S} d_i \le n^\varepsilon$.

The function that takes $x_i, d_i \mapsto x_i^{d_i}$ is a function from $\mathbb{F}_p \times \mathbb{F}_p \to \mathbb{F}_p$. In terms of the bit encodings, this takes $\{0,1\}^{(p-1)} \times \{0,1\}^{(p-1)} \to \{0,1\}^{(p-1)}$, and any such function can be realized by an NC^0 circuit of size $O(1)$. Similarly, the function that maps $d_i, x_i | i \in S \mapsto \prod_{i\in S} x_i^{d_i}$, encoded in terms of bits, maps $n^{\varepsilon}(p-1)$ bits to $(p-1)$ bits. This function can be implemented by an AC^0 circuit of size $2^{n^{\varepsilon_3}}, \varepsilon_3 < 1$. Given the value of a monomial, the value of a term, i.e., the product of the value of the monomial and its corresponding coefficient in t, can again be realized in NC^0. Finally, we sum the values of the terms by sending them to a set of appropriately modified Mod_p gates, to obtain the $p-1$ bits of the answer $t(\eta(x))$. Since the number of monomials, as explained above, is $2^{n^{\varepsilon_1}}$, the size of the circuit that computes t is $2^{n^{\varepsilon_1}+n^{\varepsilon_3}} = 2^{n^{\varepsilon'}}$ for some $\varepsilon' < 1$. (Note that the presence of Mod_p gates in $\mathsf{ACC(p)}$ is only needed at the very last stage when we sum the terms.)

- Consider the function $u(b,i) = \frac{1}{2}(1+(-1)^b \Re(\rho\omega^i))$ that maps $\mathbb{F}_2 \times \mathbb{F}_p \to \mathbb{R}$. We take n bits of all $2p$ values of $u(b,i)$ to a precision of 2^{-n} as $2pn$ bits of advice for the circuit. Then, using the circuit constructed above and comparing $t(\eta(x))$ to i for each $i \in \mathbb{F}_p$, we compute $u(b, t(\eta(x))$. We then compare this value to r as follows. If $r < u(0, t(\eta(x))$, output 1. Otherwise, output -1. Since $u(0,t(\eta(x))) + u(1,t(\eta(x))) = 1$, the value $(-1)^b$ is output with probability $u(b,t(\eta(x)))$.

Thus $\Pr_r[C(x,r) = (-1)^b] = u(b,t(\eta(x)))$. Since, for any $x \in \mathbb{F}_{p^n}^*$, $\chi(x) \in \{1,-1\}$, we find that,

$$\Pr_r[C(x,r) = \chi(x)] = \frac{1}{2}(1+\chi(x)\Re(\rho\omega^{t(\eta(x))}).$$

This is true for any $x \in \mathbb{F}_{p^n}^*$. Hence, by a standard averaging argument, there exists some setting of the random variables $r = r_0$ such that,

$$\Pr_x[C(x,r_0) = \chi(x)] \geq \frac{1}{2}(1+\frac{1}{p^n}\sum_{x\in\mathbb{F}_p^n}\chi(x)\Re(\rho\omega^{t(\eta(x))}) = \frac{1}{2}(1+|\sigma|). \tag{5.1}$$

To prove the theorem, suppose that $|\sigma| > n^{-\varepsilon}$. Then Eq. (5.1) implies that,

$$\Pr_x[C(x,r_0) = \chi(x)] > \frac{1}{2}(1+n^{-\varepsilon}).$$

This contradicts Theorem 5.2 for sufficiently small ε. Hence $|\sigma| \leq n^{-\varepsilon}$, which proves the theorem. □

Acknowledgments

The authors are grateful to Swastik Kopparty for discussions about this work and communicating some more recent results. We are grateful to the anonymous referees for careful readings of the manuscript and for comments that improved the exposition.

A Appendix: Proof of Theorem 3.1

In this section we prove the theorem:

Theorem 3.1 (Pseudorandomness of Periodic α) Let $D = \{0,1,\ldots,p-1\}$. Let t,t' be constants. Let $\sigma_0 \in D^t$ be a string of digits base p, with $\sigma_0 \neq 0^t, (p-1)^t$, and $\sigma' \in D^{t'}$.

Let $\sigma \in D^n$ be a string of base p digits of the form $\sigma'\sigma_0^{\ell}$, with $n = t' + \ell t$. Let $\alpha \in \mathbb{Z}_{p^n-1}$ be the integer whose base p representation is σ.

Then there exists $\varepsilon > 0$ such that for all sufficiently large n, depending only on σ_0, σ', we have,

$$|M_{n(p-1)/2} \cup (\alpha + M_{n(p-1)/2})| \geq (\frac{1}{2} + \varepsilon)p^n.$$

Proof. The proof, in almost all details, is identical to the one given in [3]. All differences amount to the fact that we are working in base p rather than binary; we will state explicitly where this happens. The idea is this: We must show that, if we add α to a random integer R of base-p weight $< n(p-1)/2$, then it is moderately unlikely that the sum $S = R + \alpha$ also has weight $< n(p-1)/2$. By hypothesis, the integer $\alpha \in \mathbb{Z}_{p^{n-1}}$ has a base p representation σ of period t (except for the leftmost t' digits; hence there are $\ell = \lfloor n/t \rfloor$ periods σ_0, and one string σ' left over in the t' most significant digits of α). Write R as $R'R_{\ell-1}R_{\ell-2}\ldots R_0$ and similarly S as $S'S_{\ell-1}S_{\ell-2}\ldots S_0$. While the R_i are of course pairwise independent, because of carries that arise in adding α to R, the pairs (R_i, S_i) are not independent. Nevertheless, we find that if we condition on the carries, they are independent, and this suffices for the result.

As in [3], we first analyze the distributions $R_{\ell-1}R_{\ell-2}\ldots R_0$ and $S_{\ell-1}S_{\ell-2}\ldots S_0$, and include the significant digits R', S' later.

Explicitly, we model the process of adding R and α to obtain S, with the following Markov chain. The state space is $\{(x,b)|x \in D^t, b \in \{0,1\}\}$. The transition rule is that, if we are in state (x_i, b_i), and adding $x_i + \sigma_0 + b_i$ yields the carry bit b_{i+1}, then we go into the state (x_{i+1}, b_{i+1}), where $x_{i+1} \in D^t$ is chosen at random. Begin the chain at $(x_0, b_0) = (x, 0)$, with $x \in D^t$ chosen at random, and run it for ℓ steps to obtain the sequence $(x_0,b_0),(x_1,b_1),\ldots,(x_{\ell-1},b_{\ell-1})$. Now for each $0 \leq i \leq \ell-1$, let y_i be the sum in D^t of x_i, σ_0 and b_i (omitting the carry bit). Then the pair $(x_{\ell-1}x_{\ell-2}\ldots x_0, y_{\ell-1}y_{\ell-2}\ldots y_0)$ has the same distribution as $(R_{\ell-1}R_{\ell-2}\ldots R_0, S_{\ell-1}S_{\ell-2}\ldots S_0)$.

We analyze the distribution of the sequence $b_0,\ldots,b_{\ell-1}$. This is a 2-state Markov chain determined by the stochastic matrix,

$$P = \begin{bmatrix} \frac{p^t-\alpha_0}{p^t} & \frac{\alpha_0}{p^t} \\ \frac{p^t-\alpha_0-1}{p^t} & \frac{\alpha_0+1}{p^t} \end{bmatrix} \tag{A.1}$$

where the matrix element $P_{bb'}$ ($b,b' \in \{0,1\}$) is the probability that we go from carry bit b to carry bit b'. The stationary distribution as determined by $\pi P = \pi$ can be written explicitly as $\pi(0) = 1 - \frac{\alpha_0}{p^t-1}$, $\pi(1) = \frac{\alpha_0}{p^t-1}$. Let $\pi_{bb'} = P_{bb'}\pi(b)$ be the probability that before a transition the carry is b and after it is b'.

Given any $b \in \{0,1\}$, we expect that close to $\pi(b)\ell$ of the b_i's will be b, and hence (with high probability) that about $\pi_{bb'}\ell$ transitions will be from b to b'. More precisely, fix b,b' and let z_i denote the random variable which takes on the value 1 if $b_ib_{i+1} = bb'$ and 0 otherwise. Let $Z = \sum_{i=1}^{\ell} z_i$ be the number of i such that $b_ib_{i+1} = bb'$. The mean of each z_i is $\pi_{bb'}$, and hence Z is a random variable with expectation $\pi_{bb'}\ell$. Then by the Chernoff bound, for any constant c,

$$\Pr\left[\pi_{bb'}\ell - \sqrt{\ell}(\log\ell)^c) < Z < \pi_{bb'}\ell + \sqrt{\ell}(\log\ell)^c)\right] > 1 - \ell^{-\omega(1)}. \tag{A.2}$$

We now condition the x_i on the carries $b_0,\ldots,b_{\ell-1}$. First note that each x_i is only dependent on b_i and b_{i+1} (and thus the x_i's are independent given the $b_0,\ldots,b_{\ell-1}$). In particular, x_i can be drawn randomly from D^t, under the condition that adding b_i, x_i and σ_0 produce the carry b_{i+1}. In other words, defining the distribution,

$$U_{bb'} = \{x \in D^t \,|\, \text{adding } b, x, \text{ and } \sigma_0 \text{ produces carry bit } b'\},$$

then, conditioned on $b_0, \ldots, b_{\ell-1}$, x_i is drawn uniformly at random from $U_{b_i b_{i+1}}$. As before, the distribution y_i is obtained from x_i by summing b_i, x_i and σ_0 (omitting the carry bit).

Let $\eta_{bb'}$ denote the distribution of the pair $(\mathrm{wt}_{\mathbb{F}_p}(x), \mathrm{wt}_{\mathbb{F}_p}(y))$ where x is drawn at random from $U_{bb'}$ and $y = x + \sigma_0 + b$ (omitting the carry). Let W_x and W_y denote $\mathrm{wt}_{\mathbb{F}_p}(x_{\ell-1} \ldots x_0)$ and $\mathrm{wt}_{\mathbb{F}_p}(y_{\ell-1} \ldots y_0)$, respectively, where, for each i, x_i is drawn from $U_{b_i b_{i+1}}$ (and $y_i = b + \sigma_0 + x_i$ accordingly). Then we can write

$$(W_x, W_y) = \sum_{i=0}^{\ell-1} (\mathrm{wt}_{\mathbb{F}_p}(x_i), \mathrm{wt}_{\mathbb{F}_p}(y_i)) = \sum_{i=0}^{\ell-1} \eta_i,$$

where each random variable $\eta_i = (\mathrm{wt}_{\mathbb{F}_p}(x_i), \mathrm{wt}_{\mathbb{F}_p}(y_i))$ is drawn from one of the four distributions $\eta_{bb'}$, namely $\eta_{b_i, b_{i+1}}$. Given $b_0, \ldots, b_{\ell-1}$, the η_i are independent, and by Eq. (A.2), for any choice of b, b', with high probability, close to $\pi_{bb'}\ell$ of the η_i are drawn from $\eta_{bb'}$, and are independent and identically distributed.

Hence, by the 2-dimensional central limit theorem, for each choice of b, b', for large ℓ these $\pi_{bb'}\ell$ variables are distributed according to a Gaussian with covariance matrix appropriate to $\eta_{bb'}$. Call this covariance matrix $\mathrm{COV}(\eta_{bb'})$. The pair $(W_x, W_y) = \sum_{i=0}^{\ell-1} \eta_i$ is thus the sum of four normally distributed pairs of variables, one for each choice of b, b'. The unconditioned means of W_x and W_y are both $\ell t(p-1)/2$ (since each are random strings over $D^{t\ell}$). Hence, by the additive property of the 2d normal distribution, (W_x, W_y) is normally distributed, and hence $\frac{1}{\sqrt{\ell}}(W_x - \ell t(p-1)/2, W_y - \ell t(p-1)/2)$ has mean $(0,0)$ and is drawn from a normal distribution with covariance matrix $\sum_{b,b'} \mathrm{COV}(\eta_{bb'})$.

After one more observation, we are able to exploit some general properties of the normal distribution.

We observe that W_x and W_y are not perfectly correlated. This is true because there are strings x', x'' of the same base-p weight which, when added to σ_0, yield strings of different base-p weight. We can see this as follows. Because $\sigma_0 \neq 0^t, (p-1)^t$, there is some digit of σ_0 that $= d$ where $0 < d < p-1$. Let x' have its corresponding digit $= p-d-1$ and all other digits $= 0$, and x'' have its corresponding digit $= p-d$ and all other digits $= 0$. Then $\sigma_0 + x'$ has no carry while $\sigma_0 + x''$ does. Consequently, $\mathrm{wt}_{\mathbb{F}_p}(\sigma_0 + x') = \mathrm{wt}_{\mathbb{F}_p}(\sigma_0) + p - 1$, but $\mathrm{wt}_{\mathbb{F}_p}(\sigma_0 + x'') \leq \mathrm{wt}_{\mathbb{F}_p}(\sigma_0)$ (the digit $p-d$ subtracts d and adds at most 1 to the weight via a carry).

As a consequence, defining $V_x = W_x - \ell t(p-1)/2$ and $V_y = W_y - \ell t(p-1)/2$, we find that $\Pr[V_x < 0 \wedge V_y < 0] < 1/2$. Indeed, since $\Pr[V_x < 0 \wedge V_y < 0]$ is bounded away from $1/2$, by the continuity of the probability density, there exists an ε such that $\Pr[V_x < \varepsilon \wedge V_y < \varepsilon] < 1/2 - \varepsilon$. Hence,

$$\Pr[W_x < \ell t(p-1)/2 + \varepsilon\sqrt{\ell} \wedge W_y < \ell t(p-1)/2 + \varepsilon\sqrt{\ell}] < 1/2 - \varepsilon. \tag{A.3}$$

Finally, we take the significant digits R' and S' into account, and the final carry bit c. The key point is that R' and S' both have length t', a constant, and the effect of these digits is negligible, as is also true of the effect of propagating c. Then by Eq. (A.3), we conclude,

$$\Pr[W_R < n(p-1)/2 \wedge W_S < n(p-1)/2] = \Pr[R \in M_{n(p-1)/2} \wedge S \in M_{n(p-1)/2}] < \frac{1}{2} - \varepsilon_0,$$

for some $\varepsilon_0 > 0$, which implies $|M_{n(p-1)/2} \cup (M_{n(p-1)/2} + \alpha)| > (\frac{1}{2} + \varepsilon)p^n$, for some $\varepsilon > 0$. □

References

[1] H. IWANIEC AND E. KOWALSKI: *Analytic Number Theory*. American Mathematical Society, 2004. 2

[2] S. KOPPARTY: *Algebraic methods in randomness and pseudorandomness*. Ph. D. thesis, MIT, 2010. 7

[3] S. KOPPARTY: On the complexity of powering in finite fields. In *Proc. 43rd ACM Symposium on Theory of Computing*, pp. 489–498, 2011. 1, 2, 5, 8, 9, 10, 11, 12, 14

[4] S. KOPPARTY: private communication, 2012. 2

[5] N. REINGOLD R. BEIGEL AND D. SPIELMAN: The perceptron strikes back. In *Proceedings 6th IEEE Conference on Structure in Complexity Theory*, pp. 286–291, 1991. 3

[6] A. RAZBOROV: Lower bounds for the size of circuits of bounded depth with basis $\{\oplus, \vee\}$. *Math. notes of the Academy of Science of the USSR*, 41(4):333–338. 3

[7] R. SMOLENSKY: Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proc. 19th ACM Symposium on Theory of Computing*, pp. 77–82, 1987. 3

AUTHORS

Arkadev Chattopadhyay
School of Technology and Computer Science
Tata Institute of Fundamental Research
Homi Bhabha Road, Mumbai 400005, INDIA
arkadev.c@tifr.res.in
http://www.tcs.tifr.res.in/~arkadev

Frederic Green
Department of Mathematics and Computer Science
Clark University, Worcester, MA 01610
fgreen@clarku.edu
http://mathcs.clarku.edu/~fgreen

Howard Straubing
Computer Science Department
Boston College, Chestnut Hill, MA 02467
straubin@cs.bc.edu
http://www.cs.bc.edu/~straubin